

SciDAC PDSI LANL 1st half FY07 report info

Accomplishments)

Data and Synthetic Test Code Release

LANL released 9 years of computer failure data in the spring of 2006, over 23,000 records for several thousand machines. As part of the PDSI, during this reporting period, LANL released several million usage records (job run, processors/machines used, duration, time, etc.) for the same machines that the failure data was released for. LANL also released disk failure data for a subset of these same machines. This data is publicly available and has received over 900 downloads in the last 6 months. The data has generated top paper awards at two conferences. Additionally, LANL released an important metadata synthetic benchmark the parallel file tree walk application which has received over 100 downloads. This is the first parallel file tree walker benchmark application of its kind. Additionally, LANL has cleared a process for releasing parallel I/O traces and is poised to release several gigabytes of parallel I/O traces for use in the research community.

HECURA File Systems and I/O Projects, new projects acquired, review and coordination

As part of the PDSI, LANL coordinated the first half FY07 HECIWG activities. As a part of this activity, 4 additional HECURA FSIO projects were added to the 19 existing HECURA FSIO projects. Additionally, 5 more foundations related NSF proposals were funded and added to the HECIWG FSIO umbrella for coordination, taking HECIWG FSIO university projects to 28. Additionally, a nightly BOF at USENIX FAST07 was held introducing the HECIWG FSIO and HECURA projects. At these BOF sessions, 9 of the original 19 HECURA projects were reviewed (mid year review) to ensure the HEC organizations are being well served by the research being conducted. 9 more of these mid year reviews have been scheduled.

Publications

We had no publications from LANL directly but our failure data release, synthetic application code release, and consultation assisted with the following:

Bianca Schroeder, Garth Gibson. "The computer failure data repository." Invited contribution to the *Workshop on Reliability Analysis of System Failure Data (RAF'07)* to be held at MSR Cambridge, UK. [pdf](#).

Bianca Schroeder, Garth Gibson. "Disk failures in the real world: What does an MTTF of 1,000,000 hours mean too you?" *5th Usenix Conference on File and Storage Technologies (FAST 2007)*. Winner of **best paper award**. [pdf](#)

The above paper has also been featured in an [article on slashdot](#), which so far has received more than **75,000 hits!**

Bianca Schroeder, Garth Gibson. "A large scale study of failures in high-performance-computing systems." *International Symposium on Dependable Systems and Networks (DSN 2006)*. [pdf](#)

As **one of the best DSN'06 papers** invited to *IEEE Transactions on Dependable and Secure Computing (TDSC)*.

Michael Mesnier, Matthew Wachs, Raja R. Sambasivan, Alice Zheng, Gregory R. Ganger. [Modeling the relative fitness of storage](#). International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 2007). San Diego, CA. June 12-14, 2007. ACM. **Awarded Best Paper**.

Michael Mesnier, Matthew Wachs, Raja R. Sambasivan, Julio Lopez, James Hendricks, Gregory R. Ganger, David O'Hallaron. [//TRACE: Parallel Trace Replay with Approximate Causal Events](#). Fifth Conference on File and Storage Technologies (FAST'07). San Jose, CA. February 12-13, 2007. USENIX.

Supported people

Gary Grider – LANL Staff – PDSI Task PI

James Nunez – LANL Staff

John Bent – LANL Staff

Milo Polte – CMU student summer intern at LANL

Andy Kowinski – Wisconsin student summer intern at LANL

Outreach, Collaborations and Tech Transfer

Outreach/Collaborations



On July 18, 2006, the Institute for Scalable Scientific Data Management (ISSDM) presented a tutorial on parallel I/O. This seminar was given by Dr. Robert Ross, Mathematics and Computer Science Division, Argonne National Laboratory, a renowned expert in parallel I/O programming and middleware. Rob is the co-author of the Parallel Virtual File System (PVFS) an open source file system and Message Passing Interface – I/O (parallel I/O middleware). The tutorial included both local and remote participation of approximately 35 SNL and LANL staff members and students and drew participants from at least four Groups within LANL's High Performance Computing (HPC) Division. Remote participation was via the LANL Institutes' Polycom Instructional Facility. This event showcased the Institutes' distance learning facilities which will be used to provide access to renowned experts in Computer Science (and other) fields critical to the Laboratory and to provide LANL staff with revitalization opportunities via the ability to pursue advanced degrees and continuing education in cutting edge topics.

- Collaboration with SciDAC Scientific Data Management Center –
 - Highly successful ISSDM seminar "PVFS in Production," Robert Ross, Mathematics and Computer Science Division, Argonne National Laboratory. Both local and remote participation of approx. 17 SNL and LANL staff members and students. .
- Educational Outreach:

- Infiniband, Ethernet, Computation, and Storage equipment provided to UCSC for education/research
 - Infiniband and Ethernet equipment provided to CMU for education/research
 - 2 Summer interns accepted at LANL to work on PDSI related topics
- Collaboration with Intelligent Storage Consortia
 - Planning for support for ANSI T10 OBSD open source reference
 - Planning for joint LANL ISC project in active storage
- HECIWG FSIO Coordination Activity
 - HECIWG FSIO 2007 Workshop site chosen, arrangements underway, and draft schedule published
 - HECIWG FSIO 2006 Workshop report published
 - 6 HECURA projects reviewed at FAST07. Mid year reviews for first half of HECURA FSIO projects being done to provide guidance on the research being done.
- Initial HEC POSIX Extensions published to HECEWG
- Provided support for 5 students to attend USENIX FAST07.

Tech Transfer

- Release of usage data, several million records of supercomputer usage data to match with previously released failure data was publicly released
- Release of parallel file tree walk synthetic application, this application walks a file system tree in parallel for benchmarking and other uses, this supplements the LANL MPI-IO parallel test suite.
- Initial clearing of first trace, clearing the way for future large set of trace releases

External Presentations

- Fall 2006 Topics in Computer Science Seminar Series. Local (LANL Staff) and remote (UCSC faculty and graduate students) audience participation via Polycom transmission.
 - “Devising Fault Tolerant I/O Models for Next Generation Super Computers,” John Bent (HPC-5).
 - “Scalable Server I/O Networking Architectures for Very Large-Scale Linux Clusters – PaScal I/O vs. Federated I/O,” Hsing-bung Chen (HPC-5)
 - “Failure Analysis from the Application's Point of View”, John T. Daly (HPC-4).
 - “Dynamically Probing the Linux Kernel,” Nathan DeBardleben (HPC-4)
 - “Archival Storage at LANL: Past, Present and Future,” Danny Cook (HPC-3)
- Seminar: Michael A. Bender of Stony Brook University presentation: “An Adaptive Packed-Memory Array.” This paper recently won "Best Newcomer Award" at PODS 06. (Broadcast via polycom to SNL and UCSC.)
- Standards update talk given at SC06 for HECEWG POSIX Extensions

- Educational Outreach:
 - 2 hour lecture on HPC Storage, File Systems, and I/O past, present, and future was given to UCSC graduate storage class
 - 2 hour lecture on HPC Storage, File Systems, and I/O past, present, and future was given to CMU graduate storage class

LANL SOW Tasks Status

Task Description	% of scope	% complete 1st half FY07	% complete 2nd half FY07	Milestone
1-2 LANL, SNL Participate in HPC I/O and file storage systems curriculum development at participating institute universities including courses and parenthetical degrees.	10	70		Conduct guest lectures/seminar s
1-5 ALL Develop & host I/O and file storage workshop for science application developers and users; and for I/O and file storage researchers	8	80		SC06 Apps/IO wor
1-5 LANL Sponsor SciDAC and HEC/URA/NSF I/O and file storage R&D showcase as an extension to the HEC/IWG I/O FSIO workshop to showcase the R&D	24	40		Plan and hold HEC FSIO2007 with HE status
1-5 LANL, PNNL, SNL Assist in the validation of emerging HPC storage related standards and API's such as pNFS, iSCSI Enhanced RDMA (iSER), and active storage, and enhanced POSIX I/O.	10	60		Get POSIX HECE and 1 out, Get pNF into POSIX
1-2 LANL, SNL Provided parallel I/O traces of unclassified parallel applications.	14	50		Important app IO tr public
1-2 LANL, SNL Provide parallel I/O traces of synthetic parallel benchmarks as well as source for the benchmarks to enable base lining for parallel I/O trace analysis and replay research.	14	40		Important synthetic sets public
1-2 LANL, PNNL, NERSC Collect up to a decade of supercomputer, high hp networking, and I/O and storage system reliability data, machine/environment config info, mtti/mttr) and failure cause data.	10	40		Collect/release reliability/usage/ev for systems
1-2 LANL, PNNL Collect up to a decade of supercomputer, hp networking, and I/O and file storage system usage data, including job length, size, processor usage and other usage profile data.	10	40		Collect/release reliability/usage/ev for disks (one syste